

Privacy Tailoring: A Response to the Request for Information Regarding a National Privacy Research Strategy

Contact: Privacy Research Group, Dept. of Informatics, University of California, Irvine
Kobsa@uci.edu

Summary

This document is in response to the RFI by agencies of the NITRD Program aiming to develop a joint National Privacy Research Strategy. Our reply puts forward a user-centric perspective on privacy that we feel is largely missing from the RFI. We propose to automatically *infer privacy preferences* and to *tailor privacy to each individual* as a means to reconcile the diversity of personal privacy preferences with the increasing difficulty of polling these preferences, specifically when personal data is collected by autonomous and ubiquitous/pervasive systems.

Description of the Respondent

The respondent is a research group at the Department of Informatics of the University of California, Irvine, headed by Professor Alfred Kobsa. For more than 10 years, this group has conducted empirical privacy research and has developed software architectures to safeguard privacy. The group has collaborated with a number of major IT companies on the West Coast.

Response

1. *Do we seriously want to disregard people's **individual** privacy preferences in modern computing environments?*

One of the biggest challenges of a modern National Privacy Research Strategy is the changing technological landscape. As the RFI acknowledges, these new technologies such as wearable computing, embedded computing, and cyber-physical systems “create new contexts in which privacy can be challenged.” But even more importantly, these technologies are ubiquitous, act autonomously, and often operate outside our focal awareness. This has several consequences.

First of all, these new technologies move beyond the concept of “information privacy” towards (online and real-world) behavioral tracking. Recent research has shown that a significant part of the population perceives behavioral tracking as more problematic than conventional information disclosure (Knijnenburg and Kobsa 2013). Furthermore, behavioral tracking is a *continuous* form of disclosure, which blurs the specific instances of collection and use and the subsequent opportunities for notice and control.

Second, asking users to continuously and actively address their privacy preferences fundamentally conflicts with the inherently *subliminal operation* of these new technologies.

Third, even if the user wants to specify privacy preferences, these technologies typically have no (or a severely impoverished) user interface that make the specification of such preferences very difficult or in most cases even impossible. In colloquial terms: if you think that programming the room temperature in a thermostat is difficult, try changing its privacy settings!

It is clear that new autonomous and pervasive technologies make it difficult, if not impossible, to poll individuals for their privacy decisions each and every time their personal data (informational *and* behavioral) is being processed. Full privacy control over these systems would mean that users have to make countless privacy decisions in an extremely short period of time; this would be an undue burden for them and will in many cases lead to unreasonable decisions.

The RFI addresses this problem by proposing a “responsible use framework” (4) and requesting input regarding the objectives (1), assessment (2) and broader framing (3) of this framework. The framework would thus set a baseline for the collection and use of personal information that is considered “responsible”. This approach is in line with other “soft paternalistic” policies proposed by the current administration (Nesterak 2014).

But is it even *possible* to define collectively agreed-upon limitations that would constitute “responsible use”? The RFI scratches the surface of this problematic question by acknowledging the potential existence of “diverse national/cultural perspectives on privacy”. But reality is even more complex than that: Privacy research has repeatedly shown that people have very different privacy preferences (Ackerman et al. 1999; Harris et al. 2003; Preibusch and Bonneau 2013). Moreover, people differ not just in the *degree* of information disclosure, but also in what *kinds* of information they do and do not want to disclose (Knijnenburg et al. 2013; Olson et al. 2005).

2. Privacy Tailoring

It is our belief that in a democratic and pluralistic society, there should not just be a concern for *collective* privacy; rather, people’s *individual* privacy preferences should be respected as much as reasonably possible. But how is this possible in a world where information collection and use is ubiquitous, pervasive, and subliminal? Based on preliminary efforts to solve this problem, we conclude that even if it is not possible to poll people directly for their privacy preferences, it may still be possible to *anticipate* their attitude, and then automatically *adapt* to it (Knijnenburg 2013).

For example, a number of proposals for a prediction-based approach to support users in setting their privacy preferences have been made in the field of location-sharing services (Sadeh et al. 2009; Toch, Cranshaw, et al. 2010; Xie et al. 2014). Similar work has been conducted on privacy settings for social networking services

(Ravichandran et al. 2009; Toch, Sadeh, et al. 2010). Early evaluations of “privacy recommendations” have been promising (Knijnenburg and Jin 2013), and the concept has already found its way into commercial systems such as Facebook’s “privacy dinosaur”. In the realm of wearable devices and the Internet of Things, such privacy tailoring will largely need to operate without human intervention.

The concept of “Privacy Tailoring” strikes a balance between the individually optimal level of privacy that can be attained through manual control (which is infeasible for pervasive and ubiquitous systems) and the effortless convenience of soft-paternalistic “privacy nudges” (which in their typical form constitute a one-size-fits-all approach to privacy that is arguably suboptimal for all but a few “modal” consumers). This balance results in “realistic empowerment”: it enables users to make privacy-related decisions with limited cognitive resources in an increasingly complex technological landscape.

These considerations and this early research points to a rich and diverse research agenda, specifically along the following questions:

1. What are central determinants of users’ privacy decisions? In mobile and pervasive computing environments, situational/contextual factors seem to play a far more important role than traditionally studied privacy-related attitudes and dispositions. Since those factors are often highly dynamic, their ongoing capture and determination poses additional challenges.
2. Can we identify “privacy personas”? Recent research in several areas has shown that privacy decisions are not arbitrary, but that clusters of people can be identified who exhibit similarities in the type and amount of information they chose to disclose about themselves. This points to opportunities for classifying users into privacy categories based on their characteristics and behaviors.
3. Can we learn users’ privacy preferences? Going beyond the segmentation strategies described above, can we use machine learning algorithms to predict people’s individual privacy preferences and behaviors based on their past privacy-related behaviors and situational factors?
4. Can we offer privacy-sensitive personal-information brokerage as a service? A brokerage service that is able to predict users’ privacy preferences after having learned them using (2) and/or (3) would be able to disclose a “presumably acceptable” subset of a user’s personal information to applications automatically, or grant applications a set of presumably acceptable permissions. Such a service would be extremely beneficial in ubiquitous computing and sensor environments that offer virtually no means for user interaction. A user interface to this brokerage service would in turn give back to users the possibility to view and change their inferred privacy preferences if they so desire.

Very recent work both in academia and in industry labs demonstrates initial promise of this research agenda. We strongly recommend that the agencies of the NITRD Program include automatic privacy tailoring to each individual into their

joint National Privacy Research Strategy. Privacy research has shown again and again that people exhibit stark differences in the amount and type of information they are willing to disclose. In computing environments where users can or should not be prompted for their privacy preferences all the time, automatically predicting them as far as possible will ensure that those differences are still respected rather than sacrificed to a one-size-fits-all privacy norm that will be infeasible to define.

References

- Ackerman, M. S., Cranor, L. F., and Reagle, J. 1999. "Privacy in e-commerce: examining user scenarios and privacy preferences," in *Proceedings of the 1st ACM conference on electronic commerce*, EC '99, Denver, CO: ACM Press, pp. 1–8.
- Harris, L., Westin, A. F., and associates. 2003. "Most People Are 'Privacy Pragmatists' Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits," Equifax Inc.
- Knijnenburg, B. P. 2013. "Simplifying Privacy Decisions: Towards Interactive and Adaptive Solutions," in *Proceedings of the Recsys 2013 Workshop on Human Decision Making in Recommender Systems (Decisions@ RecSys'13)*, Hong Kong, China, pp. 40–41.
- Knijnenburg, B. P., and Jin, H. 2013. "The Persuasive Effect of Privacy Recommendations," in *Twelfth Annual Workshop on HCI Research in MIS*, Milan, Italy.
- Knijnenburg, B. P., and Kobsa, A. 2013. "Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems," *ACM Transactions on Interactive Intelligent Systems* (3:3), pp. 20:1–20:23.
- Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013. "Dimensionality of information disclosure behavior," *International Journal of Human-Computer Studies* (71:12), pp. 1144–1162.
- Nesterak, E. 2014. "Head of White House 'Nudge Unit' Maya Shankar Speaks about Newly Formed Social and Behavioral Sciences Team," *The Psych Report*.
- Olson, J. S., Grudin, J., and Horvitz, E. 2005. "A study of preferences for sharing and privacy," in *CHI '05 Extended Abstracts*, Portland, OR: ACM, pp. 1985–1988.
- Preibusch, S., and Bonneau, J. 2013. "The Privacy Landscape: Product Differentiation on Data Collection," in *Economics of Information Security and Privacy III*, B. Schneier (ed.), Springer New York, pp. 263–283.
- Ravichandran, R., Benisch, M., Kelley, P., and Sadeh, N. 2009. "Capturing Social Networking Privacy Preferences:," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, I. Goldberg and M. Atallah (eds.), (Vol. 5672) Springer Berlin / Heidelberg, pp. 1–18.
- Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M., and Rao, J. 2009. "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal and Ubiquitous Computing* (13:6), pp. 401–412.
- Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L., Hong, J., and Sadeh, N. 2010. "Empirical models of privacy in location sharing," in *Proc. of the 12th ACM Intl. Conference on Ubiquitous Computing*, Copenhagen, Denmark, pp. 129–138.
- Toch, E., Sadeh, N. M., and Hong, J. 2010. "Generating default privacy policies for online social networks," in *Ext. Abstracts CHI 2010*, Atlanta, Georgia, USA: ACM, pp. 4243–4248.
- Xie, J., Knijnenburg, B. P., and Jin, H. 2014. "Location Sharing Privacy Preference: Analysis and Personalized Recommendation," in *Proceedings of the 19th International Conference on Intelligent User Interfaces*, IUI '14, Haifa, Israel: ACM, pp. 189–198.